

**SPARSE REPRESENTATIONS FOR
THREE-DIMENSIONAL RANGE DATA RESTORATION**

By

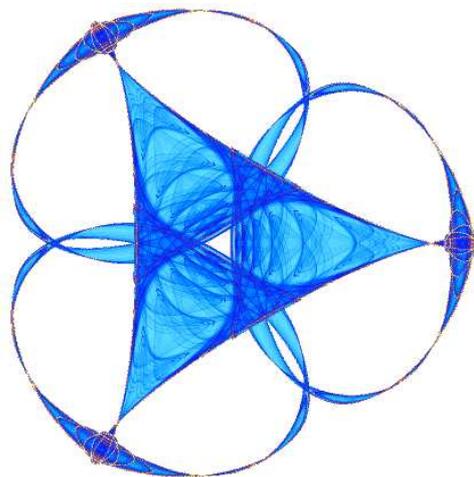
Mona Mahmoudi

and

Guillermo Sapiro

IMA Preprint Series # 2365

(February 2011)



INSTITUTE FOR MATHEMATICS AND ITS APPLICATIONS

UNIVERSITY OF MINNESOTA
400 Lind Hall
207 Church Street S.E.
Minneapolis, Minnesota 55455-0436

Phone: 612-624-6066 Fax: 612-626-7370

URL: <http://www.ima.umn.edu>

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE FEB 2011	2. REPORT TYPE	3. DATES COVERED 00-00-2011 to 00-00-2011			
4. TITLE AND SUBTITLE Sparse Representations for Three-Dimensional Range Data Restoration		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Minnesota, Institute for Mathematics and its Applications, 207 Church Street S.E., Minneapolis, MN, 55455-0436		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In this paper, the problem of denoising and occlusion restoration of 3D range data based on dictionary learning and sparse representation methods for image denoising is explored. We apply these techniques after converting the noisy 3D surface into one or more images. We present experimental results on the proposed approaches.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	Same as Report (SAR)	7	

Sparse Representations for Three-Dimensional Range Data Restoration

Mona Mahmoudi* and Guillermo Sapiro
 Department of Electrical and Computer Engineering
 University of Minnesota, Minneapolis, MN, 55455 USA
 {mahmo022,guille}@umn.edu

Abstract—In this paper, the problem of denoising and occlusion restoration of 3D range data based on dictionary learning and sparse representation methods for image denoising is explored. We apply these techniques after converting the noisy 3D surface into one or more images. We present experimental results on the proposed approaches.

Index Terms—Sparse modeling, range data denoising, occlusion restoration, range data resolution enhancement.

I. INTRODUCTION

Three-dimensional (3D) data are becoming ubiquitous. However, models obtained from low-end 3D scanners often have imperfections. For example, the raw data obtained from a low-cost 3D range scanner is usually noisy and may have occlusions and missing parts. Thus, there is an increasing need for methods for denoising and occlusion restoration of 3D surfaces in general and range data in particular. Recently, techniques based on dictionary learning for sparse representation have been widely used for image and video restoration [1], [2], [3]. In these methods, a dictionary is learned on (overlapping) patches of the image, sparsely representing those patches; that is, each patch of the image can be well approximated by a combination of only a few atoms from the learned dictionary. It has been shown that sparsely representing overlapping patches in the image with such learned dictionaries, and then combining them to reconstruct the image, results in an effective image denoising method.

In this work, we apply the framework of learned sparse representations to restore 3D surfaces, range data in particular. We also propose a new framework for filling missing parts in 3D surfaces based on ideas similar to those presented in [2]. First, we use the given 3D surface to generate one or more images, where each pixel value of an image represents the depth of a point on the 3D surface with respect to a reference plane corresponding to that image. Then, this image is denoised, using sparse modeling (dictionary learning) techniques. Finally, we regenerate the points on the 3D surface from the image by placing a point (x, y, z) on the 3D surface corresponding to each pixel (x, y) with intensity z in the image. We also introduce a method to fill the holes in the range data by applying the same sparse modeling method, by reducing the influence of the holes while estimating the representation. Our experimental results show that this technique significantly enhances the quality of the data generated by a low-cost scanner.

In many databases, the resolution of parts of the 3D range data, e.g., the low-curvature parts, are reduced for compression purposes. However, many applications require a higher or more uniform resolution for the data. Hence, an effective interpolation technique is needed to generate a higher-resolution copy of the surface, with as much similarity to the actual 3D object as possible, and without unwanted edges. As a byproduct of our proposed techniques, we increase the resolution of smoother parts of the 3D surfaces by projecting it to a high-resolution image with the same method explained above; in doing so, the details in the high-curvature areas are preserved in the image, but many empty pixels (holes) are generated corresponding to

low curvature areas without any samples. To interpolate these areas, we fill the resulting holes using the sparse modeling method, where the influence of these empty pixels are omitted in the process. We present results on a standard 3D shape database (ISDB).

In order to further improve the denoising results, in particular when dealing with structured (view-dependent) noise in the range data – which is common for structured-light low-cost scanners – we propose a *local projection method*. In this method, overlapping patches are defined on the 3D surface, and then projected to image patches for restoration, where each projection is with respect to a plane that is dependent on the corresponding 3D patch. This is in contrast to the aforementioned method – the *global projection method*, where the whole shape is first projected to a single image, before defining the overlapping patches. In the experimental results, we have compared the performance of the local projection method with that of the global projection methods, as well as a range data denoising method proposed in [4].

The remainder of this paper is organized as follows. In Section II, details about collection and preprocessing of the data, mainly from the 3D scanner, are presented. In Section III, the core algorithm for denoising 3D surfaces (range data) is presented. The method for filling the occlusions and missing information is introduced in Section IV, followed by Section V with the proposed approach for interpolation. In Section VI, the method of local projections is presented, and experimental results are provided in Section VII. Finally we conclude the paper in Section VIII.

II. DATA COLLECTION AND PREPROCESSING

In this work, we apply our proposed method to restore the range data collected from a low-cost structured-light 3D scanner [5], [6] and some range data obtained from the AIM@SHAPE shape repository,¹ and ISDB [7]. Following are details on the preprocessing step for these datasets. There was no need for preprocessing for the AIM@SHAPE dataset.

A. Range Data from Structured-Light 3D Scanner

The low-cost family of structured-light 3D scanners produce relatively noisy range data, as well as regions with missing information due to occlusions or lack of light reflection (Fig. 1). In addition, commonly used horizontal stripe patterns in the projected light add noise to the data with the shape of horizontal lines (see for example the figures in Section VII). In order to enhance the 3D data obtained from this scanner, we first convert the points on the shape to an image parallel to the camera plane. Each point with coordinates (x, y, z) corresponds to a pixel (x, y) in the image with value z , which is the distance from the camera, of the corresponding point. Having this natural and simple image representation, we apply the restoration

¹These models are provided courtesy of IMATI by the AIM@SHAPE Shape Repository.

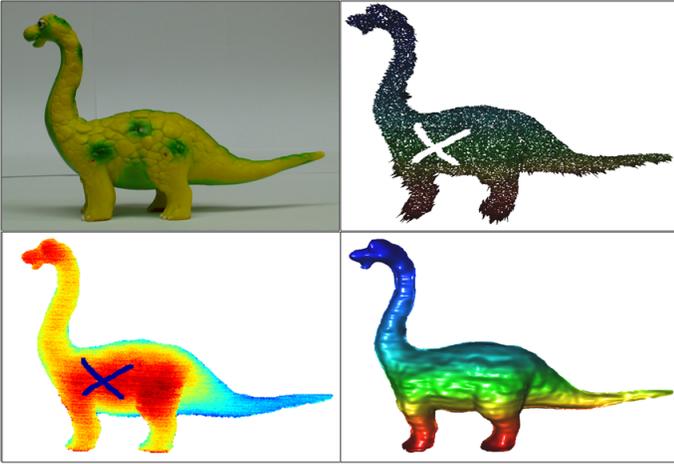


Fig. 1. *Left to right, top to bottom: a picture of the toy, the real noisy range scan of the toy with an artificial hole, noisy projected image, restored range data.*

methods explained in the following sections for denoising or filling the missing parts. Then, we show in Section VII that if we convert the restored image back to 3D points, the result will be an enhanced 3D shape.

B. 3D Shapes with Low Density

In the ISDB dataset, we have 3D shapes with very low density of vertices (sampling rate) in smooth parts and higher resolution in parts with higher curvature. Also due to very low overall resolution of shapes, our goal is to increase the density of vertices and make it uniform (see Section V for more details). Similarly to the range data, we first project the 3D shapes to image planes. For each shape, we select a plane in 3D space and project the part of the shape that is visible from such plane. We then set the pixel corresponding to each vertex on the shape as the intersection of the plane normal passing through that 3D vertex (a parallel projection). The pixel value is set to the depth of that vertex with respect to the selected image plane. Since these shapes have low densities in some parts, the higher we set the resolution of the image the more empty pixels (holes) we have in the resulting projection image. In Section V, we present a method to fill these holes in order to increase such density.

III. DENOISING SURFACES USING SPARSE MODELING

In this section, image denoising algorithms based on learned sparse representation are reviewed [2], along with the details on the specific sparse modeling techniques we use for 3D surface restoration.

Assume \mathbf{x}_0 is the clean image reshaped as a vector and \mathbf{x} is the observed noisy version of \mathbf{x}_0 . Having \mathbf{x} , we want to find the dictionary $\hat{\mathbf{D}}$ that “best” represents the patches in \mathbf{x} in a sparse fashion. In order to find $\hat{\mathbf{D}}$, the following optimization problem is addressed:

$$\begin{aligned} \min_{\mathbf{D}, \mathbf{A}} \sum_{ij} \|\mathbf{D}\alpha_{ij} - \mathbf{R}_{ij}\mathbf{x}\|_2^2, \\ \text{subject to } \|d_l\|_2^2 = 1 (l = 1..k) \text{ and } |\alpha_{ij}|_p \leq L, \end{aligned} \quad (1)$$

where L is a given constant controlling the sparsity; $p = 0, 1$ and $|\cdot|_p$ stand for the l_p norm (a pseudo-norm for $p = 0$); \mathbf{D} is the dictionary being learned, with k atoms d_l of length N ; α_{ij} is the vector of size k coefficients corresponding to the patch at location $[i, j]$, indicating the weight of each atom from \mathbf{D} in the reconstruction of the patch; and the binary matrix \mathbf{R}_{ij} extracts the

patch of size $\sqrt{N} \times \sqrt{N}$ at location $[i, j]$ from the image. The minimization is performed over the dictionary \mathbf{D} and the coding coefficients $\mathbf{A} = \{\alpha_{ij}\}$. This non-convex optimization is solved by standard techniques, altering between fixing the dictionary and performing the coding (sparse coding), and fixing the code (or active set) and updating the dictionary.

In this work, we use the unconstrained l_1 penalty,

$$\|\mathbf{D}\alpha_{ij} - \mathbf{R}_{ij}\mathbf{x}\|_2^2 + \lambda \|\alpha_{ij}\|_1, \quad (2)$$

for each $[i, j]$. For the sparse coding step, we used the LARS-Lasso algorithm [8], which is one of the most efficient algorithms in the literature for l_1 penalty problems for this type of data. We update the dictionary using a variation of the “Method of Optimal Direction” (MOD) [9], which updates the dictionary based on the current coefficients to minimize the error in Equation (1). See [10] for more details on core components of the used optimization.

In the last sparse coding step for the actual image restoration, after the dictionary has been learned, the best results were obtained imposing $\|\alpha_{ij}\|_0 \leq L$.² We applied orthogonal matching pursuit (OMP) [11], for which we can define a value ϵ as a threshold for the objective function in Equation (1), to avoid restoring the noise. Thus, when this function value in the minimization process reaches the error ϵ , which is defined based on noise variance, the OMP algorithm stops. In the case of our noisy data obtained from the low-cost camera, having large values of noise makes it harder to define a proper value for ϵ . For smaller values, a lot of noise will still be present in the restored shape, and for larger values of ϵ , we lose some of the information in the parts of shape with more details, since they need more dictionary atoms to represent them. In order to solve this problem, we modify the objective function in Equation (1),

$$\sum_{ij} \|\mathbf{D}\alpha_{ij} - (\mathbf{R}_{ij}\mathbf{x} - \text{DC}_{ij})\|_2^{2t}, \quad (3)$$

where each pixel in \mathbf{x} (after removing the DC from each patch) is taken to the power $t > 1$ which is fixed for all pixels. In this case, pixels with larger values (which are usually in higher curvature parts) have larger effect on the error than before. We have investigated the effects of this modification in Section VII. Although, this idea is limited to cases where the average shape slope is parallel to the projected image, our experimental results show that it is still helpful for very noisy data as ours. The algorithm presented in Section VI, which uses local projections with adaptive orientations, addresses this issue for shape patches with other average orientations.

Now that we have an algorithm for image denoising based on sparse modeling, we use it to denoise 3D surfaces. In order to find the 3D surface we can simply assign a point (x, y, z) to each foreground pixel (x, y) in the image whose intensity is z . The collection of these points makes the restored 3D shape.

IV. FILLING MISSING INFORMATION

Similar to images, in scanning 3D data occlusion or missing information can occur. We now investigate methods for filling/inpainting the holes in the 3D range data, assuming the location of the holes is known.³ In [2], image inpainting was investigated using sparse models. Based on this work, we address this problem for 3D range data.

The main idea is to disregard or reduce the effect of the hole pixels in the error component of Equation (2) when updating the dictionary and coding coefficients. In the dictionary update stage, we

²We thank Julien Mairal for proposing this combination and very exhaustive testing supporting it.

³These can often be easily detected from lack of signal.

remove all the patches with missing information, avoiding learning these irregular structures in the dictionary. In the last step, to find the sparse code we define a new objective function:

$$\min_{\mathbf{A}} \|\mathbf{R}_{ij}\mathbf{W} \otimes (\mathbf{D}\alpha_{ij} - \mathbf{R}_{ij}\mathbf{x})\|_2^2; \|\alpha_{ij}\|_0 \leq L, \quad (4)$$

where \mathbf{W} is an adaptive matrix of weights corresponding to each pixel, for example,

$$w(i, j) = \begin{cases} 1, & \text{known pixels,} \\ 0, & \text{hole pixels.} \end{cases} \quad (5)$$

In this part, we first subtract the DC value of each patch before estimating the coding coefficients, and add it back to the estimated patch in the reconstruction step. For the patches containing holes we set the average value of the non-hole pixels as the DC of the patch. In order to denoise the image and fill the missing information (holes), we apply Algorithm 1 to the image obtained from the damaged data. In this algorithm, we first restore all the hole pixels in the HR step, then, apply as many iterations as needed of IR and ‘‘Dictionary Update’’ to denoise the noisy shape. Having large patches in the hole filling step helps to use the information from a larger set of pixels and leads to better results. However, increasing the patch size in the restoration step makes the restored image smoother. Thus, we use large patches for HR (N_1) and smaller patches for IR (N_2).

Algorithm 1

Parameters: J (number of iterations)

Initialization: Let \mathbf{W} be the matrix of weights, which has value zero for the hole pixels and one for the rest, and let \mathbf{D}_1 ($k_1 \times N_1$) and \mathbf{D}_2 ($k_2 \times N_2$) be some learned dictionaries on one image or a dataset of images.

Hole Restoration (HR): Find the sparse code that minimize the weighted objective function (4) with learned dictionary \mathbf{D}_1 for the original image with holes, and reconstruct **only the holes**, avoiding over-smoothing in the rest of the image.

Loop: Repeat J times

- **Image Restoration (IR):** Find the sparse code that minimize the original non-weighted objective function with learned dictionary \mathbf{D}_2 for the restored image in the previous restoration step, and reconstruct the image based on this code.
 - **Dictionary Update (DU):** Update \mathbf{D}_2 based on the restored image in the IR step.
-

V. INTERPOLATION

We now apply sparse modeling methods on projected images of 3D shapes to increase the resolution of their meshes and make them more uniform. There is a lot of work in the literature on 3D compression by non-uniformly sampling 3D surfaces while minimizing the information loss. However, in some applications involving these datasets, there is a need for interpolation methods to restore a smooth and more uniform sampling of surface vertices. We propose using the sparse modeling method for image and hole restoration to both increase the 3D mesh resolution and make it uniform in 3D shapes and range scans.

First, assume we have a 3D shape or range scan with a non-uniform mesh, then we project it to an image which sufficiently preserves the information in high resolution parts and leaves some holes in the image corresponding to low resolution parts. We then apply Algorithm 1, with only one iteration of the IR step to remove the possible discontinuity in the results. In the case of noisy data, we could run more iterations of the IR and DU steps.

One of the main parameters in our proposed method is the resolution of the projection image which determines the amount of information preserved through projection. In the case of a small projection image (low resolution), more vertices will be projected

to one particular pixel, leading to information loss. However, in the case of a large projection image, there will be more possibility of holes due to lower resolution of the 3D data in the corresponding part. In Section VII, we present examples of large projection images compared to the shape which lead to a large amount of hole pixels, and also the results of interpolation of these shapes using the method proposed here.

Another idea to increase the resolution of the 3D shapes (range data) is to add empty rows and columns to the projected image (Figure 2) and fill them with Algorithm 1 again with one iteration of the IR step. Figure 3, presents an example of this method on the bunny shape.

VI. RANGE DATA DENOISING USING LOCAL PROJECTIONS

We now replace the single global projection by multiple local ones. The motivation for this idea is the existence of structured noise in range scans (local projections can also help to address restoration of full 3D objects and not just range data). Since this noise is view dependent, the averaging of restored patches with different orientations will reduce the effect of structured noise. In the previous sections, we formed the overlapping patches after (globally) projecting the surface to a single image. We propose a local method by first forming overlapping patches centered at each vertex on the 3D surface (or range data). Then, having these 3D patches, we form an image corresponding to each patch by projecting the 3D patch to the plane tangent to the surface at its center vertex. While generating the image, we keep record of the vertices projected to each pixel on the images, in order to reproduce the 3D shape after enhancement of all the projected image patches. Applying the local projection method to restore range data outperforms the global method, as demonstrated next.

VII. ADDITIONAL EXPERIMENTAL RESULTS

Our additional experimental results are presented in three parts. First, we show results of the proposed range data restoration framework for data obtained by a structured-light 3D scanner. Then, we present the results of the interpolation method on shapes from the ISDB database. Finally, we compare between the local and global projections, as well as the results presented in [4]. We present our results by 3D visualizations of the surfaces. This shows the artifacts more clearly than the depth-image visualization approaches frequently used in the literature.

The structured-light 3D scanner finds the depth of each point based on the image of the object after horizontal stripes are projected on it. Because of these stripes, an additional noise in the collected data with the shape of horizontal lines is added to the shape (Fig. 4). In some parts of the shape these lines are deeper and more difficult to remove. Also, since they exist in all the shapes, repetitive noise might be learned in the dictionary learning process. We added random holes to each image to simulate the occlusions (Fig. 4). After subtracting the DC value of the patch, we normalized all the patches by the maximum value among all the pixel values.

To obtain the results presented in this section, we have experimentally tuned the parameters. The values of these parameters are the following. In Equation (2) and Algorithm 1, the dictionary for the HR step was learned on 78 projection images. Patch size for this part was $N_1 = 225$, and the number of atoms in the dictionary was $k_1 = 550$, making it over-complete. λ was experimentally set to 0.16, $L = 12$, and the maximum iteration number in the dictionary learning (J_D) was set to 125, and best results were obtained by the following parameters in SC, $L = 2$ and $\epsilon = 2.22$. For IR, the dictionary was learned on patches of size $N_2 = 25$, and $k_2 = 63$, $\lambda = 0.48$, $L = 3$, and $J_D = 100$ for dictionary learning and $L = 2$, $\epsilon = 0.66$ for SC.

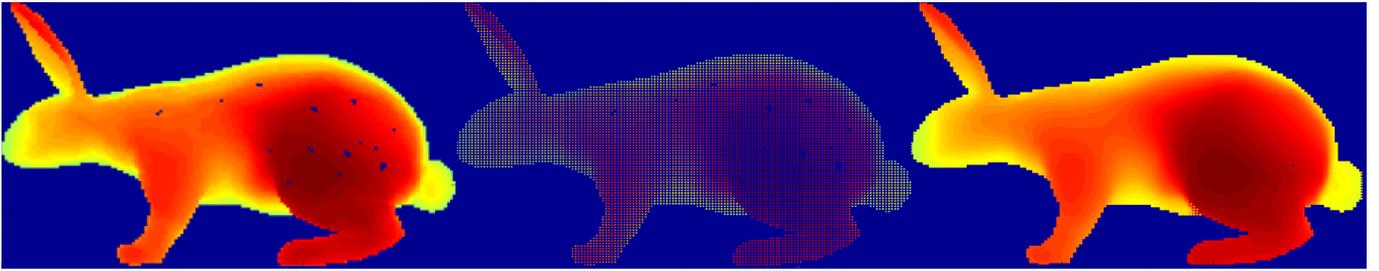


Fig. 2. Projection images in the process of interpolation by adding extra rows and columns to the image: left to right: original image (magnified low resolution image), after adding rows and columns, and the restored image.

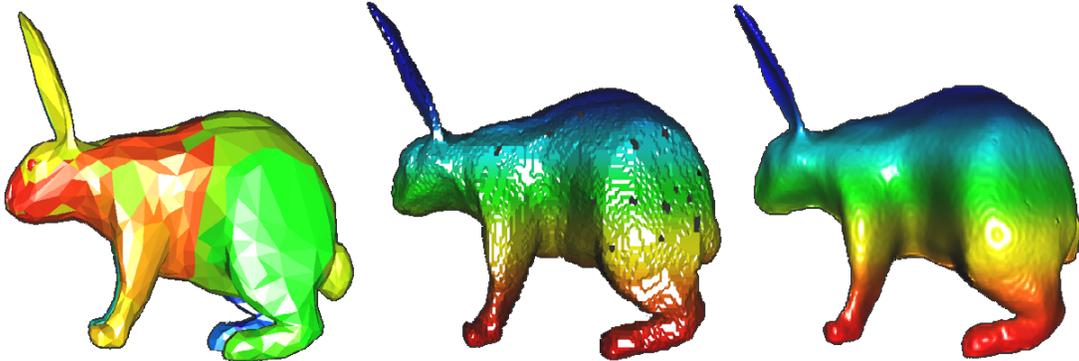


Fig. 3. The result of interpolation by adding extra rows and columns to the projection image, from left to right: original 3D shape, the result of projecting the shape to the image and back to 3D, the final result after interpolation.

In Fig. 4, examples of the scanned objects as projection images and 3D range data are presented. After the first iteration of Algorithm 1, the holes were filled and all the noise was removed except for some lines (caused by the scanning method). These residual imperfections were reduced after the second iteration, using a dictionary re-learned from the previous step restored image. The comparison between results of Equations (1) and (3) with $t = 1.5$ is also presented in Fig. 4.

Fig. 5 shows samples of learned dictionaries for different shapes. Results of applying these dictionaries to restore a pig shape are shown in Fig. 6.

The second set of experiments, on the ISDB database, investigates the effectiveness of our method for interpolation. The hole restoration parameters were the same as in the previous experiments. In the image restoration step, the only difference was the dictionary, which was learned on 78 shapes. The results of this experiments are presented in Fig. 7.

The last set of experimental results demonstrate the performance of our local projection method, described in Section VI. Moreover, we compare these results with those of our global projection method, as well as the experimental results presented in [4]. To facilitate the comparison, we have used a range data that was also studied in [4]. Due to space limitation, we only present the results for one portion of the data which contains both detailed and smooth areas. In Fig. 8 we show the original noisy range data and the results of denoising using both the global and local methods. For each of these two methods, we have used two different parameter sets, one selected to obtain a smoother outcome, and the other to preserve more details. It can be observed that the local projection method is much more successful in suppressing the view-dependent noise. This can be attributed to the fact that the structured, view-dependent noise that is present in this data will be partly hidden when the whole range data is viewed

from a single angle, while the local approach rotates each portion of the data appropriately in order to observe as much detail as possible.

For completeness, in Fig. 9 we have included the experimental results of [4] corresponding to the same portion of the range data as that in Fig. 8. Although it is difficult to quantify the qualities of the results, it can be observed that the results of our local projection method and the method of [4] are comparable for this particular sample. It is also important to note that the colors we have used to visualize the data in Fig. 8 makes the noise and details more visible than those used by [4], which we have copied in Fig. 9. This becomes apparent by comparing the left-most pictures of these two figures, both corresponding to the original noisy data.

VIII. CONCLUSIONS

In this paper, we introduced a new framework for the restoration of 3D range data. We applied sparse modeling methods on projection images obtained from 3D surfaces in order to both denoise and fill the occluded parts of the shapes. We also applied sparse modeling methods for interpolation of 3D surfaces. We also introduced a method of local projection in order to further improve the results of denoising range data. We tested these methods on data obtained from a low-cost structured-light range scanner, the ISDB database of 3D shapes, and range data from AIM@SHAPE shape repository. Our experimental results demonstrate the effectiveness of the proposed methods. A challenge for future work is to extend this work to full 3D shapes exploiting the local projection technique.

ACKNOWLEDGMENT

This work is partially supported by ARO, NGA, ONR, NSSEFF, and NSF. We thank I. Ramirez and J. Mairal for providing the code for Lasso/LARS and OMP algorithms, F. Lecumberry for his help in data collection, and E. Gordon for building the structured-light scanner and installing it in our lab.

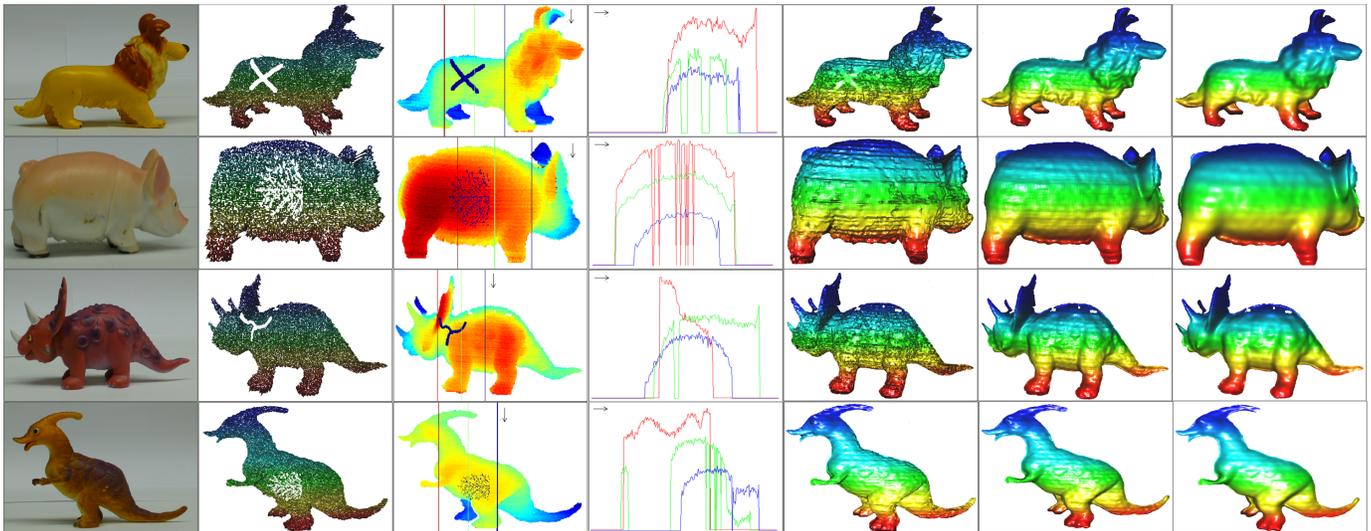


Fig. 4. Results of the proposed method in Algorithm 1 on five shapes in the dataset. From left to right, first column shows pictures of the objects, second column shows the noisy 3D shapes obtained from the 3D scanner, and third column is the converted range data image of the shapes in column two, note the horizontal lines. The fourth column shows the shifted intensity value (depth) of the pixels on the three lines shown in the images on the third column. The fifth column shows the restored shapes after the first run of Algorithm 1 with the dictionary learned on the original image and $t = 1$. The sixth column is similar to the previous column only for $t = 1.5$, and the seventh column presents the results after the second iteration of Algorithm 1 with the dictionary learned on the restored image in the first run and $t = 1.5$.

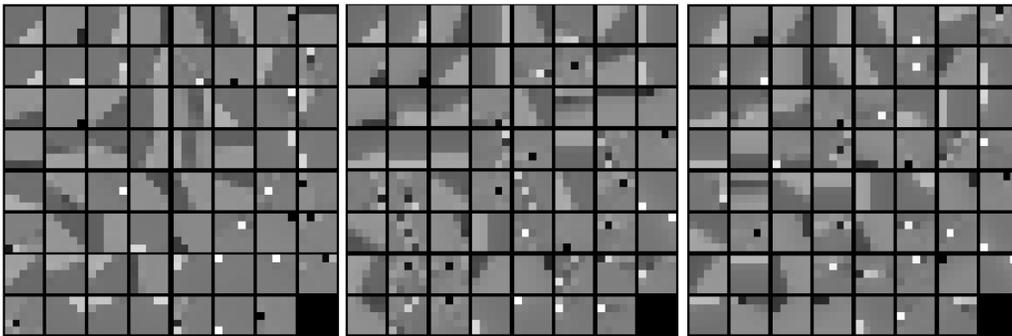


Fig. 5. Learned range data dictionaries for 78 noisy and/or occluded shapes, noisy shape of a pig with holes, and a denoised version of the pig shape, respectively. The patches are 5×5 , $k = 63$, $\lambda = 0.24$, $L = 3$, overlap of patches for the dictionary learning is 4 (full overlap), and each dictionary is learned in maximum 100 iterations.

REFERENCES

- [1] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Proc.*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [2] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Trans. Image Processing*, vol. 17, no. 1, pp. 53–69, 2008.
- [3] M. Elad and M. Protter, "Image sequence denoising via sparse and redundant representations," *IEEE Trans. Image Proc.*, vol. 18, no. 1, pp. 27–35, 2008.
- [4] O. Schall, A. Belyaev, and H. P. Seidel, "Adaptive feature-preserving non-local denoising of static and time-varying range data," *Computer-Aided Design*, vol. 40, p. 701707, 2008.
- [5] O. Rubinstein, Y. Honen, A. M. Bronstein, M. M. Bronstein, and R. Kimmel, "3d color video camera," *Proc. Workshop on 3D Digital Imaging and Modeling (3DIM)*, 2009.
- [6] A. Bronstein, M. Bronstein, E. Gordon, and R. Kimmel, "High-resolution structured light range scanner with automatic calibration," *Techn. Report CIS-2003-06, Dept. of Computer Science, Technion, Israel*, 2003.
- [7] R. Gal, A. Shamir, and D. Cohen-Or, "Pose-oblivious shape signature," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 2, pp. 261–271, 2007.
- [8] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least angle regression," *Annals of Statistics*, vol. 32, no. 2, pp. 407–451, 2004.
- [9] K. Engan, S. O. Aase, and J. H. Husoy, "Method of optimal directions for frame design," *IEEE Int'l Conf. Acoustics, Speech and Signal Processing (ICASSP)*, vol. 5, pp. 2443–2446, 1999.
- [10] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," *International Conference Machine Learning, Montreal, Canada*, June 2009.
- [11] S. Mallat and Z. Zhang, "Matching pursuit with time-frequency dictionaries," *IEEE Trans. Signal Processing*, vol. 41, pp. 3397–3415, 1993.

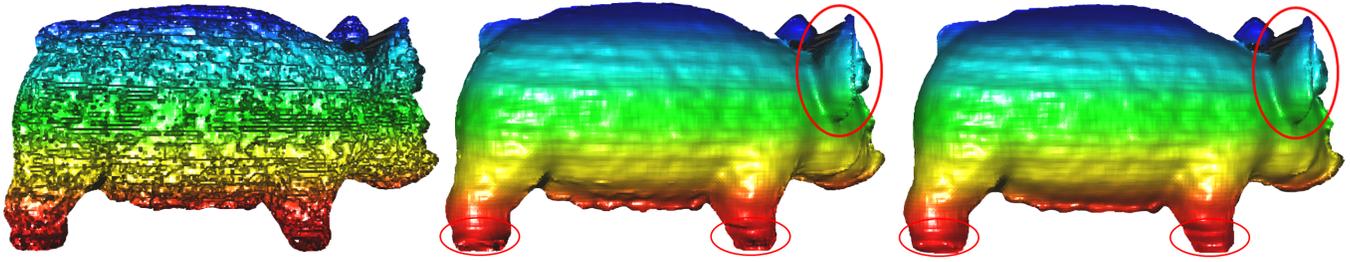


Fig. 6. Results of applying the dictionaries presented in Fig. 5 to denoise the pig shape in Fig. 4.

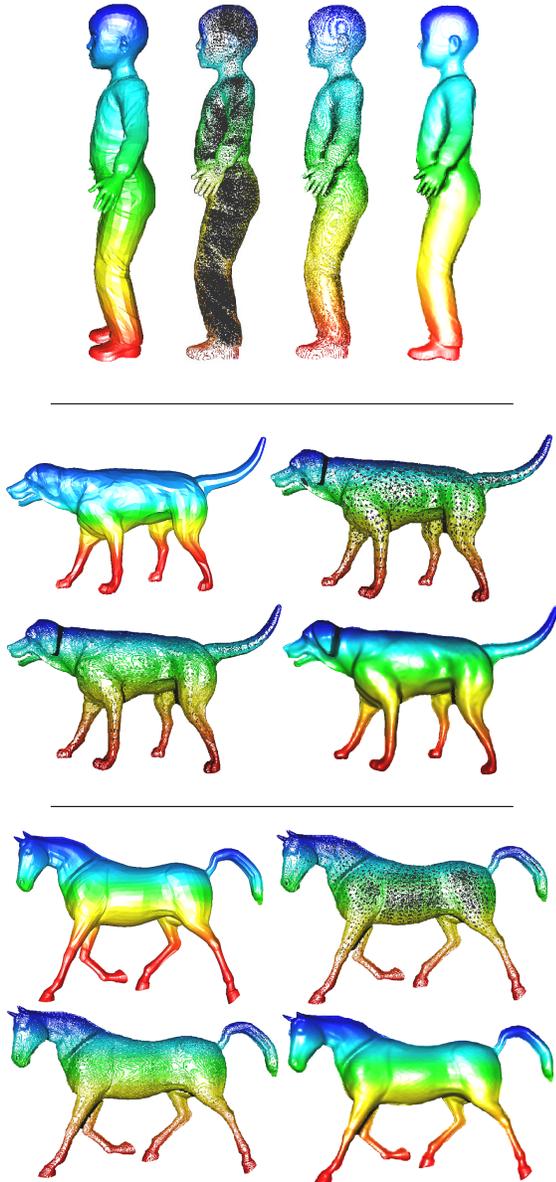


Fig. 7. Examples of interpolation method. From left to right, top to bottom: Original 3D shape, result of projecting the shape onto the image and back to 3D (dark pixels are missing information due to low density in the corresponding parts in the 3D shapes), 3D shape from the image with filled holes, final result after smoothing the previous shape.

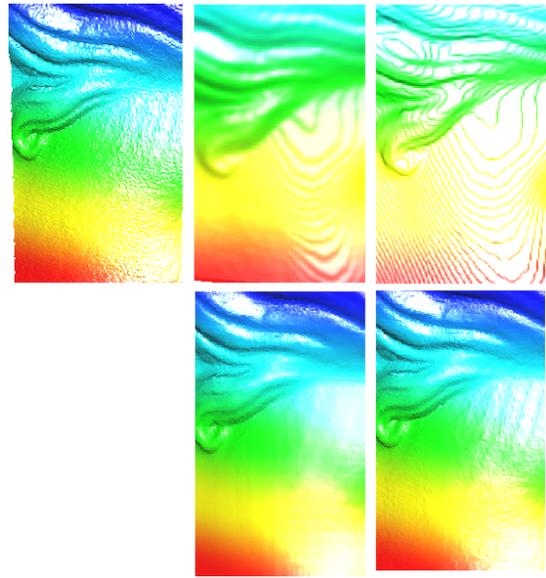


Fig. 8. Comparison of global and local projection methods: from left to right, first row: the original noisy range data and the results of applying global projection method with two different sets of values for L and ϵ in the SC step (larger L and smaller ϵ for the right shape); second row: the results of applying local projection method on the same noisy shape with similar parameters. The dictionary used for these results was learned on 78 shapes and patch size was $N = 225$.

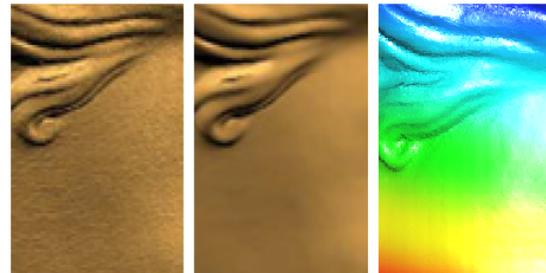


Fig. 9. The noisy shape and the result of denoising using non-local means method, reproduced from [4], and the result of the proposed local projection method.